

# A Self-adaptive Auto-scaling Method for Scientific Applications on HPC Environments and Clouds

Kiran Mantripragada  
IBM Research - Brazil  
São Paulo, Brazil  
kiran@br.ibm.com

Alécio P. D. Binotto  
IBM Research - Brazil  
São Paulo, Brazil  
abinotto@br.ibm.com

Leonardo P. Tizzei  
IBM Research - Brazil  
São Paulo, Brazil  
ltizzei@br.ibm.com

## ABSTRACT

High intensive computation applications can usually take days to months to finish an execution. During this time, it is common to have variations of the available resources when considering that such hardware is usually shared among a plurality of researchers/departments within an organization. On the other hand, High Performance Clusters can take advantage of Cloud Computing bursting techniques for the execution of applications together with on-premise resources. In order to meet deadlines, high intensive computational applications can use the Cloud to boost their performance when they are data and task parallel. This article presents an ongoing work towards the use of extended resources of an HPC execution platform together with Cloud. We propose an unified view of such heterogeneous environments and a method that monitors, predicts the application execution time, and dynamically shifts part of the domain – previously running in local HPC hardware – to be computed in the Cloud, meeting then a specific deadline. The method is exemplified along with a seismic application that, at run-time, adapts itself to move part of the processing to the Cloud (in a movement called bursting) and also auto-scales (the moved part) over cloud nodes. Our preliminary results show that there is an expected overhead for performing this movement and for synchronizing results, but the outcomes demonstrate it is an important feature for meeting deadlines in the case an on-premise cluster is overloaded or cannot provide the capacity needed for a particular project.

## Categories and Subject Descriptors

Computer systems organization [Architectures]: Distributed architectures—*Cloud computing*; Theory of computation [Design and analysis of algorithms]: Parallel algorithms—*Self-organization*.

## Keywords

Cloud computing, Self-adaptation, Load-balancing, Cloud

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the Owner/Author. Copyright is held by the owner/author(s).

ADAPT'15, Jan 21 2015, Amsterdam, Netherlands.  
Part of ADAPT Workshop proceedings, 2015 (arXiv:1412.2347).  
ADAPT/2015/01.

bursting, Dynamic federation to cloud.

## 1. INTRODUCTION

High Performance Computing (HPC) platforms are designed to be explored by intensive computing applications to obtain performance, such as seismic analysis (for the exploration of oil and gas and natural disasters), weather and flooding forecasts, computational fluid dynamics, DNA sequencing, simulations of electromagnetic equation, among others. They are dedicated to provide their maximum throughput; and both applications and resources are usually fine-tuned to perform as optimally as possible for their specific challenges. Companies, government laboratories, and research/educational institutions acquire such HPC resources built on clusters of supercomputers, which are typically expensive to acquire and maintain. However, there is a considerable amount of time that such environments can stay idle, while other times it can be overloaded or even has capacity limitations that, together, prevents a particular project or simulation deadline to be achieved.

Cloud computing, on the other hand, proposes a resource-shared model in which users allocate resources on demand from providers only when necessary in a “pay-as-you-run” model, spending budget only for utilized processing time. Cloud was initially created to serve mostly Web applications deployed on Virtual Machines (VM) that share the same physical hardware with other VMs at the same time. As it matured, this model was coined Public Cloud and another model has emerged to deal with private data and also HPC requirements: the Private Cloud. It focuses on hosting computing intensive industrial and scientific workloads that deal with sensitive and large amounts of data – on demand, but in a private and specialized space. This latter model have then specific requirements that are challenging to be met by Cloud providers, like the use of specialized resources (e.g., accelerators, GPUs, Infiniband, parallel file systems, etc), transfer and storage of large data sets, visualization, security and privacy, control, and a collaboration infrastructure. Currently, providers are starting to offer Private Clouds with such specialized resources and this movement can eventually allows a hybrid execution, i.e., the use of on-premise HPC Clusters together with Clouds for boosting the execution platform at a certain period and to allow particular projects to meet their deadlines.

We propose a method to partition tasks and data of an application in both environments in an unified view. We use a seismic application as use case. This particular application can reflect the characteristics of a broad range of

scientific and industrial applications: it is computational demand, data and task parallel, communication bound, utilizes solvers for differential equations, produces large amounts of data to be visualized as post-processing, etc. Our approach is able to monitor the application at runtime and to predict its total execution time so that it can perform a self-adaptation towards a Cloud Bursting, i.e., automatically shifts part of the data from a Cluster to be processed by the Cloud, in an unified and synchronized way, with the goal to meet a time deadline.

This article is organized as the following: Section 2 describes the proposed approach by emphasizing how the method performs the self-adaptation, auto-scaling, and automatic bursting, i.e., the application adapts its execution on both platforms at runtime; Section 3 exemplifies the approach using a seismic application as an example and shows preliminary outcomes with a discussion; and Section 4 finalizes the proposal with related work, followed by conclusions on Section 5.

## 2. DYNAMIC AND SELF-ADAPTIVE CLOUD BURSTING

Let us consider an application that is data and task parallel, and which also evolves over time. Initially, the user starts the application according to specific parameters: *data* (domain size), *tasks*, number of *timesteps*, *partitions* of the data to be executed on-premise and in the cloud (data can be initialized fully placed on-premise). The system immediately starts to monitor the execution time for each timestep. The simulation continues to execute until the system detects that the *threshold* time, i.e., the *deadline* will not be satisfied. This situation can happen for a number of reasons, like concurrency in the local cluster, nodes down, dynamic change of the deadline, less on-premise capacity for a given problem, etc. Once the "time monitor" detects a change in the estimated execution time, the system verifies this new estimation against the deadline time provided by the user. It is important to note that initial sequential timesteps are monitored in order to reason whether they are predictable. In most cases, these industrial and scientific applications are simulations that evolves over time by solving a number of partial differential equations represented as matrices operations. Analyzing these characteristics, the timesteps can have similar execution times, but working on distinct data and, thus, being fairly predictable.

If the monitoring module detects that the system will not attend the time threshold, this estimated time is used to compute the number of cores it would be needed to fit the simulation time within the deadline. This is done using a pre-processing phase to estimate the behavior of the application or it could also be an input parameter given by the user. Then, the system automatically starts the re-partitioning phase to migrate part of the domain to be computed by the elastic Cloud platform. Figure 1 and the following steps depict the proposed solution:

1. Monitoring the application and analyzing estimated execution time to decide about a Cloud bursting;
2. If a bursting is needed, save current state;
3. Compute the number of cores to be allocated in the Cloud extended environment;




Figure 1: Steps of the self-adaptive method

4. Compute the size of domain (part of the full domain) to be placed in the Cloud;
5. Move current state data to new nodes (Cloud-burst);
6. Assimilate current state as initial conditions;
7. Restart simulation at the stopped step with new configuration (on-premise cluster and Cloud);
8. Synchronize the simulation at each timestep running in both environments and merge results.

In order to compute the estimated deadline and for validation purposes, we empirically defined the computational behavior of the application. By executing it for several cores and nodes configuration, a function that can represent the overall behavior was estimated:

$$L_{cloud}(c) = -A \ln c + B \quad (1)$$

$$L_{cluster}(c) = -D \ln c + E. \quad (2)$$

Equations 1 and 2 represent the application response as a function of number of cores, respectively for the Cloud environment and for an on-premise Cluster platform. The coefficients A, B, D and E were empirically computed for the experimental application using a small test job (small data set and few timesteps) as pre-processing—see Section 3. The number of cores needed to attend the time threshold would satisfy only the Cluster environment, but since there is a possibility to place part of the domain to be executed in a Cloud environment, one must consider the network delays for migrating the data and restart the application as the reduced computing performance inherent in a Cloud environment when compared to the cluster. The correction factor for the difference between the performances (Cluster x Cloud) can be further defined as follows:

$$K = \frac{L_{cloud}}{L_{cluster}}$$


$$K = \frac{-A \ln c + B}{-D \ln c + E}$$

where  $K$  is the computational performance correction factor and A, B, C and D are coefficients computed for a given application (Section 3).

Once the system computed the new number of cores required to attend the deadline, the correction factor can be applied:

$$c_n = (c - c_{cluster}) * K \quad (3)$$

where  $c_n$  is the number of cores in the extended environment,  $c$  is the estimated number of cores (Equation 2) and  $c_{cluster}$  is number of cores available in the on-premise cluster.



**Figure 2: Domain split for Cloud-Bursting ( $\gamma$  is the size to be placed in a Cloud)**

The next step is to split the domain into partitions to be executed in the hybrid environment. In order to simplify

**Table 1: Hardware configuration of the Cluster and Cloud environments**

	Cluster	Cloud
processor (GHz)	2.80	2.60
cores per processor	10	4
cache size (KB)	25600	20480
total memory (KB)	132127868	65711672
network	ethernet/ infiniband	ethernet

the method, one of the two dimensions in domain (width and height) was fixed. As depicted in Figure 2, we chose to fix the height dimension to compute the value of  $\gamma$  as a function of execution time. That approach allows to define a linear relationship between the execution time ( $t$ ) and the domain size through the value of  $\gamma$ :

$$f(\gamma) = t = a\gamma + b \quad (4)$$

where the linear coefficient  $a$  and the offset  $b$  must be computed for each application.  $\gamma$  must be any integer value, since it represents the number of column partitions to be placed in the extended environment.

The Equation 4 can be written as a function of time ( $t$ ):

$$f^{-1} = g(t) = \gamma = \frac{t - b}{a}. \quad (5)$$

The difference between estimated total time and the time for deadline is applied in Equation 5 to compute the size of the partial domain ( $\gamma$ ) to be executed by the Cloud environment.

### 3. EMPIRICAL STUDY

The goal of this study is to analyze the approach previously described for the purpose of creating a proof-of-concept model based on the application characteristics. Such analysis was carried out from the perspective of the developers of cloud bursting scientific applications targeting a seismic application. Obviously, the study involves two computing environments: on-premise Cluster environment and the Cloud computing environment. Both environments have different hardware (*e.g.*, memory, CPU) and software configurations (*e.g.*, operating system). The cloud environment is the *IBM SoftLayer platform*<sup>1</sup>. Table 1 shows the hardware configuration of both environments. We have picked the cloud node configuration offer that was the most similar to the one in our local cluster for more correct evaluations.


#### 3.1 Target application: FWI

Full Waveform Inversion (FWI) is a numerically challenging technique based on full wavefield modeling of a geological domain that extracts relevant quantitative information from seismograms [7]. When dealing with realistic physics-based elastic partial differential equations formulation and accurate discretization techniques, such as high order Spectral Element Method, the forward modeling becomes particularly challenging from the computational point of view. This type of modeling often requires grid representation of

<sup>1</sup><http://www.softlayer.com/>

the order of billions of elements with interpolation polynomials of at least  $4^{th}$  degree in each spatial dimension, which entails hundreds of billions of variables computed for the forward solver. Furthermore, in the course of local optimization, the aforementioned process or its adjoint are repeated multiple times per non-linear inversion iteration. The computational complexity associated with FWI applications demands tremendous efforts to set up and maintain a distributed scientific HPC infrastructure. These applications are also embedded in complex business models with various stakeholders. Therefore, having FWI applications offered as a service in the Cloud can bring several benefits to their users.

Summarizing, millions of shots in different positions (i.e., independent tasks) are calculated over the same domain (i.e., data) and each shot solution is composed of several traces, evolving along timesteps. At the end, the generated images are composed to produce the final outcome. This process is depicted in Figure 3. Above, multiple traces are produced by for shot by the acquisition ship that is moving over an area of exploration. The ship carries the compressed air gun along with an array of seismograms that capture wave reflections produced by different subsurface geological layers. Below, a produced timestep, in different point-of-views, is visualized. It represents a partial result of the reconstructed wave using our implementation of the FWI over a small region of interest of the area that contains a specific salt dome (we used the SEG/EAGE salt dome model).



**Figure 3: Examples of FWI data collection and remote visualization of the forward solver**

This particular implementation (C++) is parallelized by means of OpenMPI (over data and tasks) and uses the Eigen3<sup>2</sup> library for matrix and vector manipulation, as well as ParaView<sup>3</sup> to provide (remote) visualization. The application was executed several times in each environment and Table 2 presents the details of the executions.

**Table 2: Execution details**

number of elements (x axis)	600
number of elements (y axis)	600
polynomial order (x axis)	4
polynomial order (y axis)	4
degrees of freedom	5764801
domain size (x axis)	9000
domain size (y axis)	6000
number of timesteps	3000

### 3.2 Empirical Formulations

To validate our assumptions, we have performed some experiments on the target application. The first equations were empirically determined through the results acquired from the simulations. The chart on Figure 4 depicts that both curves have similar behavior, but as the number of processor-cores decreases to 10, the elapsed time for the FWI simulation showed an increase of 150% ( $10^{5.4}/10^{5.0}$ ) worse than the on-premise cluster environment. On the other hand, as the number of processor-cores increases up to 40, the elapsed time of *SoftLayer* environment was only around 50% ( $10^{4.3}/10^{4.1}$ ) longer than the cluster environment, tending to be the same with the increasing of cores. Furthermore, it is worth mentioning that these preliminary tests were executed with only a small number of cores and nodes to verify the behavior of Cloud, while running a high CPU and I/O intensive application.

From the Figure 4, we defined the equation 1 and 2 as follows:

$$L_{cloud}(c) = -0.77 \ln c + 7.1 \quad (6)$$

$$L_{cluster}(c) = -0.65 \ln c + 6.5 \quad (7)$$


$$g(t) = \gamma = \frac{t - 231.18}{7.46} \quad (8)$$

By solving equations 6 to 8 for  $c$ , we can dynamically compute the value of  $c_n$  in equation 3 and  $\gamma$  in 8. Those values are, respectively, the number of cores and the size of the domain placed in the Cloud environment. After applying the movement, the system is now capable to execute the simulation and satisfy the deadlines. This process is repeated while the system detects that the deadline time limit would not be attended for any reason, thus delaying the execution.

### 3.3 Discussion

<sup>2</sup><http://eigen.tuxfamily.org/>

<sup>3</sup><http://www.paraview.org/>




**Figure 4: Execution time as a function of the number of processors**

In this research, we propose a method that allows to dynamically move part of a domain to Cloud environments. In addition, the benefit of running the numerical solver in a Hybrid infrastructure (cluster and a private Cloud) showed interesting behavior when we forced the local HPC cluster to exhaust its limited resources. Besides application performance gains, the proposed self-adaptive method reduces the time span needed for the “time to production” infrastructure, since there is no need to wait the acquisition, installation, and setup processes as would be with in acquiring new local resources, i.e., the co-existence in a balanced way of both platforms can bring better price-performance to such applications.

We performed the simulations over a few number of configurations and load-balancing strategies. The graph provided by Figure 5 shows the application response for several grid sizes while holding one dimension constant. Such an experiment allowed to understand how the domain split strategy would work when migrating part of the domain to an external environment. This study – execution time when varying processor cores – allowed to understand how to provision the Cloud cores for a given application size and deadline. We then observed that the inverse process would also be an interesting way to understand how the execution time can be extrapolated to a specific number of cores. The self-adaptive strategy is based on the equations 1 to 2, which allows to determine a few parameters to burst the application towards the Cloud. Those parameters, *viz.*  $\gamma$  and  $c_n$ , are sufficient to define the external infrastructure.

The overhead of the monitoring and partitioning components can be neglected. Results show that total message size is only 21KBytes size. The second-level partitioning strategy is based on stripes, which reduces the amount of communications between partitions and between environments. It is important to note that the chosen load-balancing method over cores inside an environment is the simple greedy-based algorithm, where the striped-partitions are assigned to processor cores as long as they are available in each node. This choice reduces the inter-nodal communications as the neigh-



**Figure 5: Empirical results for execution time (seconds) as a function of  $\gamma$**

bors partitions tend to be assigned to only one cluster node. In our case study, shots are independent tasks and traces have dependencies inside a shot.

On the other hand, the overhead caused by saving the actual state (checkpoint), the transferring corresponding data to the Cloud, and the provision the Cloud nodes need to be accounted. These values are not neglected and need to be considerable shorter than letting the application to finish its execution on the local cluster. Although we need to carefully analyze these measurements, in a seismic application – which usually takes months to produce the final result – such automatic *checkpoint-restart* process is cost-effective as the total execution time of this application is inferred in the beginning of execution. We understand that this assumption would cause an inaccurate estimation and we are extending this work to measure and include such overhead as an offset value in equations 1 and 2.

Despite of the potential feature for the system to be elastic, in the sense that infrastructure can expand and shrink as needed, the current approach just computes and applies expansions on the Cloud environment. Further studies and advances on this system can even consider to reduce the amount of processors when the deadline threshold can be achieved with reduced infrastructure.

Finally, the proposed method can be incorporated in a framework to be reused in different applications. For that, one needs to investigate the applications behavior before specifying the coefficient values as described in section 2. Considering that we simplified the mesh partitioning by fixing one of the dimensions, we are planing to extend this work by investigating the full 2D mesh partitioning and also full 3D numerical solvers.

## 4. RELATED WORK

High Performance Computing applications are being tested on Cloud platforms. Works like [2], [3] and [4] performed a performance evaluation of a set of benchmarks and complex HPC applications on a range of platforms, from supercomputers to clusters, both in-house and in the cloud. These studies show that a Cloud can be effective for such appli-

cations mainly in the case of complementing supercomputers using models such as cloud burst and application-aware mapping to achieve significant cost benefits. Although these findings do not propose an automatic and adaptive approach for using both environments, their empirical studies opened the opportunity for proposals of tools that promote a hybrid approach based on these environments, like ours.

Analyzing Cloud as stand alone execution platform for HPC applications, like seismic, the authors of [5] evaluated the Linpack workload on the Amazon EC2 cloud. Their conclusions indicate that the tested cloud environment has a potential, but it is not mature to provide a price-performance for HPC applications. [6] also evaluated the EC2 for a number of kernels used by HPC applications, coming also to a conclusion that such cloud services need an order of magnitude in performance improvement to better serve the scientific community. It is hard to evaluate one provider or another, but they are evolving to offers that are private and with specialized infrastructures, like with GPUs and Infini-band. In our study, we evaluated the *SoftLayer* Cloud (virtualized nodes) and preliminary results indicate the environment as cost-effective in budget and performance when at least combined with on-premise clusters in a dynamic changing scenario.

More recently, [1] evaluated a computational fluid dynamics application over a heterogeneous environment of a cluster and the EC2 cloud. The results indicated that there is a need to adjust the CPU power (configuration) and workload by means of load-balancing. We are in line with this study and went further with the present work – a dynamic self-adaptive method for application load-balancing over a hybrid platform composed of cluster and cloud.

## 5. CONCLUSION

We presented a first step towards a framework for self-adaptation of industrial and scientific applications in terms of being executed on a hybrid and heterogeneous environment composed of on-premise HPC clusters and the Cloud. It dynamically monitors and reasons when to migrate part of the computation from one platform to the other at runtime to maximize performance and meet deadline constraints. We demonstrated the core method applied to a seismic application, which is data and task parallel and could reflect the behavior of a broad range of industrial and scientific applications. The method is based on a self-adaptation of the application at runtime to shift part of the computation to the Cloud when the prediction of the total execution time is identified to overcome the given execution deadline (which could also change dynamically). The approach prevents additional expensive capital expenditure and has an intrinsic overhead, but due to the on-demand cloud elasticity of nodes provisioning, the needed amount of nodes can be aggregated to achieve the deadline.

Ongoing work is based on the tool refinement, including overheads detection and shaping into a framework to be incorporated on the IBM's *SoftLayer* Private and Public Clouds and/or even on the *Bluemix platform*<sup>4</sup> as a service targeting FWI as a service. Moreover, the roadmap includes to work with data from a real seismic exploration field acquisition and therefore a more robust cluster capable of processing the problem. Also, we intend to proceed with the

analysis of nodes composed by CPU and GPU, thus including another level of heterogeneity inside the execution nodes (i.e., dynamically scheduling over the processors, like CPU and GPU, as reported in [8]).

Finally, the approach presented here at first sight might seem like a counter-intuitive idea: replace a straightforward embarrassingly parallel implementation of a challenging real world problem with a more sophisticated, self-adaptive, auto-scaled, synchronized, communication intense, domain partitioned approach. However, we have shown that a seamless combination of on-premise Clusters with the Cloud has an important contribution to boost the performance of computing intensive HPC applications with better price-performance and time-to-production.

## 6. ACKNOWLEDGMENTS

This work has been partially supported by FINEP/MCTI under grant no. 03.14.0062.00.

## 7. REFERENCES

- [1] M. B. Belgacem and B. Chopard. A hybrid HPC/cloud distributed infrastructure: Coupling EC2 cloud resources with HPC clusters to run large tightly coupled multiscale applications. *Future Generation Computer Systems*, 42(0):11 – 21, 2015.
- [2] A. Gupta, L. V. Kale, F. Gioachin, V. March, C. H. Suen, B.-S. Lee, P. Faraboschi, R. Kaufmann, and D. Milojicic. The who, what, why and how of high performance computing applications in the cloud. In *Proceedings of the IEEE International Conference on Cloud Computing Technology and Science (CloudCom)*, 2013.
- [3] A. Gupta and D. Milojicic. Evaluation of hpc applications on cloud. In *Open Cirrus Summit (OCS)*, 2011 Sixth, 2011.
- [4] G. Mateescu, W. Gentzsch, and C. J. Ribbens. Hybrid Computing—Where HPC meets grid and Cloud Computing. *Future Generation Computer Systems*, 27(5), 2011.
- [5] J. Napper and P. Bientinesi. Can cloud computing reach the top500? In *Proceedings of the combined workshops on UnConventional high performance computing workshop plus memory access workshop*, 2009.
- [6] S. Ostermann, A. Iosup, N. Yigitbasi, R. Prodan, T. Fahringer, and D. Epema. A performance analysis of ec2 cloud computing services for scientific computing. In *International Conference on Cloud Computing (CloudComp'09)*. 2010.
- [7] J. Virieux and S. Operto. An overview of full-waveform inversion in exploration geophysics. *Geophysics*, 74(6):WCC1–WCC26, 2009.
- [8] A. Binotto, M. Wehrmeister, A. Kuijper and C. Pereira. Sm@rtConfig: A context-aware runtime and tuning system using an aspect-oriented approach for data intensive engineering applications. *Control Engineering Practice*, 21(2):204–217, 2013.

<sup>4</sup><https://www.bluemix.net>